

NOTES ON FIRING THEORY

Alan R. Washburn

Naval Postgraduate School
Monterey, California

January 2000

DEMO QUALITY UNRESTRAINED 4

20001020 081

INTERNET DOCUMENT INFORMATION FORM

A. Report Title: Notes on Firing Theory

B. DATE Report downloaded From the Internet: October 13, 2000

Report's Point of Contact: (Name, Organization, Address, Office Symbol, & Ph #): Naval Postgraduate School, Monterey, CA

D. Currently Applicable Classification Level: Unclassified

E. Distribution Statement A: Approved for Public Release

F. The foregoing information was compiled and provided by:

DTIC-OCA Initials: LL **Preparation Date** October 19, 2000

The foregoing information should exactly correspond to the Title, Report Number, and the Date on the accompanying report document. If there are mismatches, or other questions, contact the above OCA Representative for resolution.

1. INTRODUCTION	1
2. SINGLE SHOT KILL PROBABILITY	2
2.1 Definitions	2
2.2 Cookie cutter weapons	3
2.3 Carleton weapons	4
2.4 Other damage functions	5
3. MULTIPLE SHOT KILL PROBABILITY	6
3.1 Simultaneous independent shots	6
3.2 Simultaneous dependent shots	9
3.3 Area targets/Multiple error sources	15
3.4 Sequential shots with feedback	17
4. DEFENSE OF ONE TARGET	20
4.1 Known attack size	20
4.2 Bayesian defense	24
4.3 The maximum cost defense	26
4.4 Prim-Read defense	27

1. Introduction

These notes are intended to be tutorial in nature, rather than comprehensive. The reader who desires a comprehensive treatment should see [1], which contains additional references to the considerable literature that exists on “coverage problems.” It seems to be the nature of the subject that there are a great many conceptually similar cases and sub-cases, each requiring a different mathematical treatment. Our goal here is to describe and summarize the main ideas, recording in the process only those results for which simple expressions are available.

The material in the first three sections of these notes is devoted to computations of the probability of “killing” a target with possibly several “weapons”, with the effectiveness of each weapon depending on a two-dimensional miss distance. The same mathematics applies to computations of such things as the probability of “detecting” a target with “sensors”; the only essential feature is that the crucial event must either happen or not. Partial damage is not permitted — each shot either kills the target or leaves it unscathed. This assumption is often not realistic, but it nonetheless must serve because practically all analysis is based on it.

When a density function for firing errors is required, it will invariably be taken to be bivariate normal. The Central Limit Theorem is the justification for this assumption, since a firing error can usually be thought of as being composed of several more or less independent parts. This is not to say that all firing errors are normal, but the normal distribution is nonetheless a natural benchmark.

The reader should already be aware that these notes will only be easily digestible to someone whose background in probability includes the idea of bivariate density functions. A knowledge of differential and integral calculus will also be assumed. It will be useful to have the Firing.xls Excel workbook available.

2. Single Shot Kill Probability

2.1 Definitions

The basic interaction between weapon and target is through the “damage function” $D(r)$, which is defined to be the probability that the target is killed by a weapon if the relative distance between them (the miss distance) is r . Determination of the damage function is in practice done through some combination of theory and experiment; we will invariably assume the function to be known, avoiding any implied statistical estimation problems. Note the implicit assumption of radial symmetry of damage effects, since the damage function does not have an angular argument.

The damage function can be thought of as a conditional kill probability. The kill probability P_K is obtained by averaging over the miss distance. Let $f(x,y)$ be the bivariate density of the position of the target relative to the weapon. Then, since $r = \sqrt{x^2 + y^2}$,

$$(2-1) \quad P_K = \iint D(\sqrt{x^2 + y^2}) f(x,y) dx dy,$$

where the lack of limits means that the integral is to be taken over the whole plane. Sections 2.2 through 2.4 deal with various special cases of (2-1).

If the target were uniformly distributed within some large area A , then (2-1) would be (substituting $f(x,y) = 1/A$),

$$(2-2) \quad P_K = \frac{1}{A} \iint_A D(\sqrt{x^2 + y^2}) dx dy,$$

where the notation indicates that the integral is now taken only over the area A . However, since A is by assumption large, (2-2) is approximately the same as $P_K = a/A$, where

$$(2-3) \quad a = \iint D(\sqrt{x^2 + y^2}) dx dy, \text{ or}$$

$$(2-4) \quad a = 2\pi \int_0^\infty r D(r) dr.$$

Formula (2-4) was obtained from (2-3) by introducing polar coordinates. The quantity “ a ” is the “lethal area” of the weapon, and serves as a scalar measure of weapon size. It plays a role in coverage problems that is similar to the role of sweep width in Search Theory, but it has dimensions of area, rather than length.

Although it is not logically necessary, the damage function is typically non-increasing. As long as this is true, it is sometimes convenient to imagine that each weapon has a random “lethal radius” R associated with it, and that any target within R of the weapon will be killed. Recalling the meaning of $D(r)$, it must evidently be the case that

$$(2-5) \quad D(r) = P(R > r).$$

If $D(r)$ is differentiable, one can go further and discover the probability density function of R :

$$(2-6) \quad f_R(r) = -\frac{d}{dr} D(r).$$

The area covered by the weapon is πR^2 , so it should come as no surprise that $a = \pi E(R^2)$, where $E(\)$ denotes expectation; demonstration of this is left as exercise 3.

2.2 ***Cookie cutter weapons***

The conceptually simplest kind of weapon is one for which the lethal radius R is a constant, in which case the lethal area is of course πR^2 . If the firing errors are circular normal (by which we mean that the standard deviation of the error in all directions is the same number σ) and centered on the target, then the two-dimensional density function of the error is $f(x,y) = \exp(-\frac{1}{2}(x^2 + y^2)/\sigma^2)/(2\pi\sigma^2)$, and (2-1) reduces to

$$(2-7) \quad P_K = 1 - \exp\left(-\frac{1}{2}R^2/\sigma^2\right).$$

Unfortunately, most departures from the above assumption about errors result in much more complicated expressions for P_K . If the circular normal error distribution is offset

from the target, for example, or if the normal distribution is not circular, then evaluation of (2-1) must be done by numerical integration or some other numerical technique. The fact that tables of the results are available has not prevented the application of (2-7) in circumstances where it is at best approximate, for example to problems where the down-range error is significantly larger than the cross-range error.

Formula (2-7) is sometimes expressed in the form

$$(2-8) \quad P_K = 1 - (.5)^{(R^2/\text{CEP}^2)}$$

where the CEP or “circular error probable” is by definition the radius of the circle that contains half the firing errors. For a circular normal distribution, CEP is related to σ by $\text{CEP} = \sigma\sqrt{2 \ln 2} = 1.1774\sigma$.

2.3 Carleton weapons

This section could also have been titled “diffuse Gaussian weapons”, since both terms are used in practice. The damage function has the form $D(r) = \exp(-\frac{1}{2}r^2/b^2)$ for some scale factor b . The lethal area of such a weapon is $2\pi b^2$. Figure 1 compares $D(r)$ for a




Figure 1

Carleton weapon with two other damage functions that will be discussed in the next section. The Carleton weapon is evidently “sloppier” than the others. Whether this feature makes the Carleton assumption more realistic depends on the damage mechanism. Weapons that kill by fragmentation generally have a sloppier damage function than those that kill by overpressure.

The Carleton assumption combines very nicely with the assumption of normal errors to produce a simple, general expression for P_K . If the center of the error distribution is (μ_X, μ_Y) , and if the standard deviations of the X and Y errors are (σ_X, σ_Y) , then (2-1) can be evaluated analytically:

$$(2-9) \quad P_K = \frac{b^2}{\sqrt{(b^2 + \sigma_X^2)(b^2 + \sigma_Y^2)}} \exp\left(-\frac{1}{2}\left(\frac{\mu_X^2}{b^2 + \sigma_X^2} + \frac{\mu_Y^2}{b^2 + \sigma_Y^2}\right)\right).$$

In the special case where $\mu_X = \mu_Y = 0$ and $\sigma_X = \sigma_Y = \sigma$, (2-9) reduces to

$$(2-10) \quad P_K = b^2/(b^2 + \sigma^2),$$

which is comparable to (2-7). There is no cookie cutter counterpart to (2-9). While it is true that the cookie cutter weapon is conceptually simpler than the Carleton, it is equally true that the Carleton is analytically simpler than the cookie cutter.

2.4 Other damage functions

It was pointed out in Section 2.1 that any non-increasing damage function can be interpreted as the probability law for a random lethal radius R . The Carleton damage function, for example, has associated with it the density function $f_R(r) = (r/b^2) \exp(-\frac{1}{2} r^2/b^2)$, which is a Rayleigh density. It is perhaps more natural to deal with the random variable R^2 , since R^2 is directly related to area covered; for the Carleton damage function, R^2 is an exponential random variable with mean $2b^2$.

It is possible, of course, to reverse the process: begin with some convenient density for R or R^2 and then discover the associated damage function by integration. One

convenient class of damage functions (the Gamma class) can be obtained by assuming that $\frac{1}{2}R^2/b^2$ has the Gamma density $[(\alpha x)^{\alpha-1}/\Gamma(\alpha)]\exp(-\alpha x)$ for some $\alpha > 0$, in which case the Carleton damage function is the special case $\alpha = 1$, the cookie cutter is obtained in the limit as $\alpha \rightarrow \infty$, and $E(R^2) = 2b^2$ for all α ; i.e., every member of the class has the same lethal area $2\pi b^2$. The associated damage function is

$$(2-11) \quad D_\alpha(r) = 1 - P\left(\alpha, \frac{\alpha r^2}{2b^2}\right)$$

where $P(\alpha, x)$ is the incomplete Gamma function (in Excel, $P(\alpha, x)$ is GAMMADIST($x, \alpha, 1, \text{true}$) — see Figure 1 or Firing.xls for plots of $D_\alpha(r)$ versus $\frac{r}{\sqrt{2b}}$).

The Gamma class is convenient because it has both scaling (b) and shaping (α) parameters, and also because there is a simple formula for P_K when the firing error is circular normal with standard deviation σ and centered on the target:

$$(2-12) \quad P_K = 1 - \left(1 + \frac{b^2}{\alpha\sigma^2}\right)^{-\alpha}; \quad \alpha > 0.$$

(2-10) is the special case $\alpha = 1$, and (2-7) is the limiting case as $\alpha \rightarrow \infty$.

Another class of density functions for R^2 with both a shape and a scaling parameter is the class of lognormal densities. There turns out to be little to recommend this class in terms of analytic convenience; there are no counterparts to (2-11) and (2-12), for example. Nonetheless, the class is widely used to model the effects of nuclear weapons [2].

3. Multiple Shot Kill Probability

3.1 Simultaneous independent shots

Suppose that n shots are fired at a target, and let q_i be the probability that the i^{th} shot fails to kill the target. The numbers q_i may be obtained from one of the formulas in

Section 2 or by some other method. Since all shots are by assumption independent, the probability that all n miss the target is the product of the miss probabilities, so

$$(3-1) \quad P_K = 1 - q_1 q_2 \dots q_n.$$

Formula (3-1) takes on a particularly simple form if the shots are cookie cutter and the firing errors are circular normal centered on the target. Let R_i and σ_i be the lethal radius and error standard deviation of the i^{th} shot. Then $q_i = \exp\left(-\frac{1}{2} R_i^2 / \sigma_i^2\right)$ from (2-7), and

therefore

$$(3-2) \quad P_K = 1 - \exp(-X/2), \text{ where}$$

$X = R_1^2 / \sigma_1^2 + \dots + R_n^2 / \sigma_n^2$. The quantity X can be thought of as a measure of the effectiveness of an arsenal of weapons against a particular target. The target dependence can be eliminated if lethal radius scales in a known manner with the energy yield Y of the weapon. If the kill mechanism is overpressure, for example, then $R_i = KY_i^{1/3}$ where K is a target dependent constant, and therefore $X = K^2 [Y_1^{2/3} / \sigma_1^2 + \dots + Y_n^{2/3} / \sigma_n^2]$. The quantity in [] is a target independent measure of effectiveness for the group of weapons taken as a whole. It differs from “counter military potential” (CMP) only in the scale factor required to convert standard deviation to circular error probable (CEP) for circular normal weapons (see Sec. 2.2). The CMP of a group of weapons is

$$(3-3) \quad \text{CMP} \equiv Y_1^{2/3} / \text{CEP}_1^2 + \dots + Y_n^{2/3} / \text{CEP}_n^2.$$

Counter Military Potential is one of several quantities that have been used to compare arsenals of nuclear weapons, with yield being measured in megatons and CEP in nautical miles. Note that CMP is very sensitive to accuracy; doubling all yields increases CMP by the factor $2^{2/3} = 1.6$, whereas halving all CEP's increases CMP by the larger factor $2^2 = 4$. In the 1970's, this fact was sometimes used to make the point that the small (relatively) but accurate nuclear arsenal of the United States was actually more potent than the large

but inaccurate arsenal of the Soviet Union. Tshipis [3], for example, estimated in 1974 that CMP was 22000 for the US and 4000 for the SU. An alternative measure of effectiveness for an arsenal is “equivalent megatons” (EMT), according to which the Soviet Union had the larger arsenal during the same period. The definition of EMT is

$$(3-4) \quad \text{EMT} \equiv Y_1^{2/3} + \dots + Y_n^{2/3}.$$

Since $Y_i^{2/3}$ is proportional to R_i^2 , EMT is essentially a measure of the total lethal area of the arsenal. Whether EMT or CMP is the more appropriate measure is discussed further in Section 3.3. The EMTCMP page of Firing.xls compares the 1978 ICBM arsenals of the USA and the USSR using both measures.

If a total of C units of CMP are applied to a target, then the kill probability is of course still a function of the hardness of the target. For nuclear weapons making overpressure kills, with hardness h being measured in pounds per square inch, an approximate formula is

$$(3-5) \quad P_K = 1 - \exp(-7.51Ch^{-75}).$$

For example, a one megaton weapon with a CEP of .25 nautical miles will kill a 1000 psi target with probability $1 - \exp(-(7.51)(16)(.0056)) = .49$. Sixteen such weapons would be equally effective if the CEP were 1 nautical mile.

Figure 2 shows an application of the EMT idea to various historical ships. $E8RPM$ is the equivalent rate of applying 8" rounds per minute, plotted against range for each of five ships. Since yield (Y) is proportional to weight and weight to the cube of dimension, $Y^{2/3}$ is proportional to dimension squared. Thus one 16" round is equivalent to four 8" rounds, etc. Reducing all firepower to a common scale permits each ship to have a single “weight of broadside” versus range curve, independent of the target. Aircraft carriers can be portrayed on the same scale by calculating the rate at which aircraft can deliver 500-pound bombs, each of which is equivalent to two 8" rounds.

Figure 2 does not tell the whole story, of course, since no information about accuracy or armor is given. Figure 2 does make it clear that the Enterprise and (say) the New Jersey were very different weapon systems, and that the passage of 35 years produced an aircraft carrier much more powerful than the Enterprise.




Figure 2. Equivalent 8" Rounds Per Minute For Five Ships

3.2 Simultaneous dependent shots

The firing errors dealt with in the previous section were *dispersion errors*, by which is meant that the weapon impact points relative to the target are a collection of independent random variables. In this section we assume the additional presence of a *bias error*, by which is meant a normally distributed error that is common to all shots. This error might be due to a misalignment between the aiming and launching systems, to an

error in target location, or to any other effect(s) that introduces an error component common to all shots. The result is frequently as illustrated on the cover; the impact points relative to the target are tightly grouped (indicating small dispersion errors) but in the wrong place. One can think of the bias error as being the center of gravity of the group, and as the dispersion errors as being deviations from the center of gravity. We shall use the notation that (σ_U, σ_V) are the (horizontal and vertical, say) standard deviations of the bias error, whereas the independent dispersion error for each shot has standard deviations (σ_X, σ_Y) .

It is no longer possible to proceed by first finding the single shot kill probability and then invoking an independence assumption to obtain a simple expression for P_K , since the independence assumption is falsified by the bias error. We will find, in fact, that there are no simple exact expressions for P_K in any circumstances. The primary reason for this is that the shots should in general be aimed in some sort of a pattern, rather than directly at the target, which means that P_K should now be “the probability that the target is killed when the shots are aimed in an optimal pattern”. The implied optimization problem is non-linear in the variables ($2n$ of them if there are n weapons), and with no special structure that can be exploited. The best that can be hoped for in such circumstances (other than solutions to specific problems that are important enough to justify the work involved in evaluating a large number of patterns) is some rules of thumb that take the form of approximations. In deriving these approximations, it will be convenient to imagine that the only source of bias is an error in target location, but the approximations are valid regardless of the source of bias or even if there are several sources (see Sec. 3.3).

Our first approximation to P_K is an upper bound obtained by making two unrealistic assumptions that are clearly favorable to the marksman. One assumption is that there are no dispersion errors, and the other assumption is that the marksman can exchange his

weapons for any other weapon or weapons with the same total lethal area. If $\sigma_U = \sigma_V = \sigma$, it is clear that the marksman would always prefer to have a single large cookie cutter weapon that he would aim directly at the target, or more precisely at the mean location of the target. If the total lethal area of n weapons is na , then the lethal radius of such a single weapon would be $R = \sqrt{na/\pi}$, and the resulting kill probability would be (from (2-7)) $1 - \exp\left(-\frac{1}{2}R^2/\sigma^2\right) = 1 - \exp(-(na)/(2\pi\sigma^2))$. More generally, the best weapon for our privileged marksman to choose is a cookie cutter with the same elliptical shape as the iso-probability contours of the error distribution, and the resulting bound is

$$(3-6) \quad P_K \leq 1 - \exp(-z), \text{ where } z = \frac{na}{2\pi\sigma_U\sigma_V}.$$

Formula (3-6) was obtained by essentially assuming away all the overlap that is caused by dispersion errors, circle-packing problems, and (effectively) non-cookie cutter weapons. The expression $1 - \exp(-z)$ should therefore be expected to be an accurate approximation in circumstances where overlap is expected to be a minor problem. Seven circles, for example, pack rather nicely into one circle without very much overlap. Figure 3, which is taken from the Patterns sheet of Firing.xls, shows the upper bound lying well above two other approximations that are introduced below.

A different kind of approximation is based on the idea that overlap is inevitable, and that one should expect the amount of overlap to be whatever happens “at random”. More precisely, the total lethal area na is assumed to be *in effect* so much confetti, with the marksman being able to control the density of confetti on a large scale, but not the small-scale tendency of the flakes to overlap one another. Now, if d square inches of confetti are scattered on a one-inch square, or in other words if the density or coverage ratio is d , then the fraction of the square that remains uncovered is $\exp(-d)$ as long as the flakes are sufficiently small (see exercise 6). The conditional kill probability is therefore

$1 - \exp(-d)$, and the marksman's problem is to determine d in such a manner that the (unconditional) kill probability is maximized.




Figure 3

Assume that $\sigma_U = \sigma_V = \sigma$, and that the marksman scatters the confetti uniformly over a circle with radius r in the hope that some flake covers the target. This is the strategically uniform, locally random (SULR) case. Within the circle, the coverage ratio is $d = na/\pi r^2$, so the probability of killing the target given that the target lies within the circle is $1 - \exp(-d)$. The probability that the target is actually in the circle is (from (2-7)) $1 - \exp(-\frac{1}{2}r^2/\sigma^2)$, so the unconditional kill probability is

$$(3-7) \quad p(r) \equiv \left[1 - \exp\left(-\frac{1}{2}r^2/\sigma^2\right)\right] \left[1 - \exp(-na/\pi r^2)\right].$$

Note that the first factor in (3-7) is 0 if $r = 0$, whereas the second is 0 if $r = \infty$. There must be a maximizing value for r . The value is $r^* = \sigma(4z)^{1/4}$, where $z = na/2\pi\sigma^2$, as can be verified by showing that $(d/dr)p(r^*) = 0$. Upon substituting r^* into (3-7), one obtains the SULR formula

$$(3-8) \quad P_K = p(r^*) = (1 - \exp(-\sqrt{z}))^2.$$

Formula (3-8) also holds when $\sigma_U \neq \sigma_V$, provided that $z = na/2\pi\sigma_U\sigma_V$ and that the confetti is scattered uniformly over an optimally sized ellipse. Figure 3 shows that the SULR formula provides a much smaller estimate of P_K than does (3-6).

The final approximation is the same confetti approximation except that the coverage ratio can be any function $d(x, y)$ of two spatial coordinates, subject of course to being non-negative and to the constraint that the total amount of confetti used must be na . This includes the case where $d(x, y)$ is constant within some region and 0 outside it, so we should expect the current approximation to be larger than (3-8). This is the strategically optimal, locally random (SOLR) case. Formally, the optimization problem is:

$$\text{maximize } \iint f(x, y)[1 - \exp(-d(x, y))] dx dy$$

$$\text{subject to } d(x, y) \geq 0 \text{ for all } x, y$$

$$\text{and } \iint d(x, y) dx dy = na,$$

where $f(x, y)$ is the bivariate normal density function with standard deviations (σ_U, σ_V) . The solution can be found in [4], together with a discussion of how the optimal coverage ratio $d^*(x, y)$ can be used as a guide in designing effective patterns. The optimal function $d^*(x, y)$ is

$$(3-9) \quad d^*(x, y) = \frac{1}{2} \left(\sqrt{8z} - \frac{x^2}{\sigma_U^2} - \frac{y^2}{\sigma_V^2} \right)^+,$$

where the + indicates that $d^*(x, y)$ is to be 0 rather than negative, and where $z = na/2\pi\sigma_U\sigma_V$ as usual. Note that the confetti should be most dense at the origin, with the density falling off gradually to 0 on the $(8z)^{1/4}$ standard deviation ellipse. Outside of this ellipse there should be no confetti at all. The result of substituting $d^*(x, y)$ into the objective function is the SOLR formula

$$(3-10) \quad P_K = 1 - (1 + \sqrt{2z}) \exp(-\sqrt{2z}),$$

which is usually identified as the “ $\sqrt{-}$ formula”, even though (3-8) is equally deserving of the name. The $\sqrt{-}$ formula is also shown in Figure 3. There is not much difference between (3-8) and (3-10). Once the total lethal area has been conceptually reduced to confetti, it turns out not to be crucial that its distribution be exactly (3-9). The $\sqrt{-}$ formula is much more widely used as an approximation than (3-8).

An example may be of some help at this point. Suppose that there are four cookie cutter weapons, with $R = 7.5$, and that the error standard deviations are $\sigma_U = \sigma_V = 7.5$, $\sigma_X = \sigma_Y = 1$. By exhaustive trial and error computations, it can be determined that the exact best pattern is a square of side 11.7, and that the associated kill probability is .80. Since $z = 4\pi(7.5)^2/2\pi(7.5)^2 = 2$, the three approximations are .865, .594 (SOLR), and .573 (SULR). The upper bound is considerably closer to the truth than either of the confetti approximations. The confetti approximations can be made to look better by letting the weapons be Carleton with the same lethal area, in which case the approximations don't change but exact computations reveal that the best P_K is only .69, achieved by aiming the four weapons in a square of side 10. If the dispersion error is in addition increased from 1 to 5, the approximations still don't change, but the best possible P_K decreases to .62.

Since neither σ_X , σ_Y , nor any feature of the damage function other than lethal area enters the computation of z , it is clear that one could find cases where the actual kill probability is even smaller than the confetti approximations. In fact, one has only to consider any problem where the shots are nearly independent, since $z = \infty$ when σ_U or σ_V is 0. In problems where the bias errors dominate the dispersion errors, however, the confetti approximations can usually be thought of as lower bounds on P_K .

Given all the above considerations, we offer the following procedure for obtaining an approximate P_K in the general case where both bias and dispersion are present:

- (a) If dispersion dominates bias, determine the “equivalent” dispersion standard deviations $\sigma'_X = \sqrt{\sigma_X^2 + \sigma_U^2}$ and $\sigma'_Y = \sqrt{\sigma_Y^2 + \sigma_V^2}$, solve the single shot kill probability problem, and then use (3-1) to obtain an approximate P_K .
- (b) If bias dominates dispersion, and if the “packing problem” can probably be solved without much overlap (nearly cookie-cutter weapons, dispersion small compared to lethal radius as well as bias, etc.), use (3-6).
- (c) If bias dominates dispersion, and if it is clear that the best pattern will involve substantial overlap, use one of the confetti approximations.

The above rules are not exhaustive, since there are certainly cases where neither type of error dominates the other, and in any case the resulting estimate of P_K is only an approximation. An accurate P_K can only be obtained by evaluating (by Monte Carlo simulation, for example — see exercise 8) sufficiently many patterns to be sure of having discovered the best one.

3.3 Area targets/Multiple error sources

Section 3.2 is often applicable even when there are multiple sources of error. Suppose, for example, that

- (a) the location of a target relative to some known datum is E_1 .
- (b) all shots are to be fired from a platform whose location relative to the same datum is E_2 .
- (c) each shot has an individual firing error E_3 due to trembling on the part of the marksman.
- (d) an additional firing error is introduced due to an unknown wind velocity E_4 .
- (e) E_1, E_2, E_3 , and E_4 are all independent, normal random variables with 0 mean and variances $\sigma_i^2, i = 1, 2, 3, 4$.

It is necessary to classify each of the four errors as either “bias” or “dispersion.” E_1 and E_2 are clearly bias, since the positions of the target and the platform are the same for each

shot. E_3 is clearly dispersion, since each shot has an independent dispersion error that is different from all the rest. E_4 might be bias if the unpredictable part of wind velocity were constant in space over the length of time required to fire the shots (the predictable part is irrelevant, since the marksman could allow for it in aiming), or it might be dispersion if the wind were very gusty. Assume that wind error is actually dispersion. Then, making the natural assumption that the four error types are independent of each other, and noting that it is only the total bias and the total dispersion that affect the fate of the target, the equivalent bias and dispersion variances are $\sigma_1^2 + \sigma_2^2$ and $\sigma_3^2 + \sigma_4^2$, respectively, and Section 3.2 can be applied to the equivalent errors. The principle being used is the theorem that the variance of a sum of independent random variables is the sum of the variances.

It is remarkably easy to handle area targets within this scheme. Suppose that E_1 only applies to the center of the target, about which point the value density (value per unit area) of the target is $V(x, y)$, and that the meaning of P_K is “the average fraction of the target value killed”. If the total target value is V_0 , then $V(x, y)/V_0$ has all the properties of a density function, and can in fact be interpreted as the density function of the location E_0 of a randomly selected “test element” of the target. With this interpretation, P_K is “the probability of killing the test element (a point target)”, and E_0 is a bias error. In other words, any area target can be handled by converting the value density of the area target to an equivalent density function of a bias error, and then proceeding as if the target were a point target. This is especially easy to do, of course, if $V(x, y)/V_0$ turns out to be bivariate normal. Suppose, for example, that $V(x, y)/V_0$ is circular normal with standard deviation $\sigma_0 = 80$ ft., that E_1, E_2, E_3 , and E_4 are all circular normal with standard deviations 10, 20, 30, and 40 ft. respectively. Assuming as before that the wind error is dispersion, the equivalent dispersion is $\sigma_X = \sigma_Y = \sqrt{30^2 + 40^2} = 50$ ft., and the equivalent bias is $\sigma_U = \sigma_V = \sqrt{10^2 + 20^2 + 80^2} = 83$ ft. One could now proceed as in Section 3.2, probably by

ignoring the dispersion error and using the $\sqrt{-}$ formula to estimate P_K , which is now interpreted as the maximum possible expected fraction of the target killed by an optimal pattern.

The fact that area targets introduce an effective bias error is important in determining whether CMP or EMT is a better measure of effectiveness for an arsenal of weapons (see Sec. 3.1). Since (3-3) was derived under the assumption that the only firing error was dispersion, we can say that CMP is the proper measure if the targets are point targets and if the bias errors are very small. If the effective bias (including the effects of target size) dominates the effective dispersion, however, then EMT is more appropriate. Thus (to conclude the comparison that was begun in Sec. 3.1), the United States nuclear arsenal in the 1970's was more effective against well located, hard targets such as ICBM silos, but the Soviet Union arsenal was more effective against cities, which are well located area targets, or against submarines, which are poorly located point targets. Dispersion is almost irrelevant for either of these latter target types, even though it is crucial for the former.

Some firing errors are neither bias nor dispersion, but instead vary by a small amount between shots. Wind, for example, may fall in this category, as may aim point wander in rapid-fire weapon systems. Firing problems associated with such errors are difficult.

3.4 Sequential shots with feedback

Sections 3.1, 3.2, and 3.3 all deal with firing problems where no information feedback is available between shots. Such feedback can be quite valuable in terms of resources required to kill the target. The purpose of this section is to examine two firing procedures that take advantage of it: shoot-adjust-shoot (SAS) and shoot-look-shoot (SLS).

In the (one-dimensional) SAS procedure, it is assumed that an observer provides a signed miss distance X_i after the i^{th} shot. These observer reports are useful because they help the marksman to estimate whatever bias error B is present, and thereby to adjust his

i^{th} aim point A_i to take account of it. Assuming that the dispersion error is E_i for the i^{th} shot, the fundamental relationship is

$$(3-11) \quad X_i = B + E_i - A_i; \quad i > 1.$$

The aim point A_{i+1} can be determined by the marksman from the observed miss distances X_1, \dots, X_i , and should in all cases be the marksman's best estimate of the unknown bias B . Since $B + E_i$ is an inaccurate but unbiased observation of B , and since $B + E_i = X_i - A_i$, the minimum variance estimate of B after i shots is:

$$(3-12) \quad A_{i+1} = \frac{1}{i} \sum_{j=1}^i (X_j + A_j) = B + \frac{1}{i} \sum_{j=1}^i E_j; \quad i > 1,$$

and therefore, with all aim points but the first being given by (3-12),

$$(3-13) \quad X_{i+1} = E_{i+1} - \frac{1}{i} \sum_{j=1}^i E_j; \quad i > 1.$$

Assuming that the dispersion errors are normal, independent, identically distributed random variables with mean 0 and variance σ^2 , $E(X_{i+1}) = 0$ and (from (3-13))

$$(3-14) \quad \text{Var}(X_{i+1}) = \sigma^2 + \sigma^2/i = \sigma^2(i+1)/i; \quad i > 1.$$

Formula (3-14) applies to every shot except the first, which we regard as a "calibration shot" ($A_1 = 0$) that is incapable of killing the target, with subsequent shots being "for effect". Alternatively, the bias error can be regarded as being unknown but so large that the chances of success for the first shot are negligible. In either case, the desired effect can be obtained by taking (3-14) to hold for $i = 0$, in which case $\text{Var}(X_0) = \infty$. Since the miss distances after the first are all independent of each other, the probability of kill with a fixed number of shots can be obtained with the same independence argument that leads to (3-1).

Equation (3-12) can be rearranged to look like

$$(3-15) \quad A_{i+1} - A_i = X_i/i; \quad i \geq 1.$$

Equation (3-15) states that the aim point for the next shot should be corrected by a decreasingly small fraction of the previous miss distance. In this form it is sometimes called "Whistler's rule."

Suppose now that the SAS procedure is carried out independently in each of two dimensions, using n cookie cutter shots with lethal radius R , including the calibration round. The two dimensional miss distances will then be circular normal with variance given by (3-14), and therefore, using (2-7) and (3-1) in the same manner as in Sec. 3.1,

$$(3-16) \quad P_K = 1 - \exp\left(-\frac{R^2}{2\sigma^2}\left(\frac{1}{2} + \frac{2}{3} + \dots + \frac{n-1}{n}\right)\right).$$

Note that the effectiveness (CMP – see Sec. 3.1) of the i^{th} shot, compared to its effectiveness in a problem with no bias error, is $(i-1)/i$. The SAS procedure is evidently not completely effective in getting rid of the effects of the bias error, except in the limit when there are many shots. There is nonetheless a reasonable sense in which it is the optimal aim adjustment procedure.

The shoot-look-shoot (SLS) procedure involves feedback about whether the target has been killed, rather than about miss distance. The advantage of such information is that it helps prevent the assignment of additional weapons to a target that has already been killed. In the extreme case where the number of looks is unbounded, the marksman can even adopt the strategy "fire until the target has been killed", in which case the problem is not to compute P_K (which is 1.0), but rather to investigate the random variable $N \equiv$ "number of shots required to kill the target". If, for example, the shots are all independent with kill probability p , then N is a geometric random variable with mean $1/p$. More generally, if q_i is the miss probability of the i^{th} in a sequence of independent shots, then

$$(3-17) \quad E(N) = \sum_{n=0}^{\infty} P(N > n) = 1 + q_1 + q_1 q_2 + q_1 q_2 q_3 + \dots$$

There is not a great deal more that can be said about the SLS procedure as applied to a single target. SLS is more naturally applied to problems with several targets, as in Sec. 4.1 below.

In general, knowledge of miss distance is not sufficient to determine whether the target is killed, so there are firing problems where SLS is present but not SAS, as well as vice versa, or both may be present. When both SLS and SAS are present, one can consider the problem of computing $E(N)$ for a given aim adjustment procedure, or even the problem of determining the procedure that minimizes $E(N)$. Computation of $E(N)$ for the aiming procedure (3-12) is left as exercise 10.

4. Defense of One Target

4.1 Known attack size

Assume that each of n attackers will kill its target with probability p if not intercepted, and that the defender has m interceptors, each of which will kill an attacker with probability ρ , and all of which are to be used against the n attackers. The defender's goal is to maximize the probability that the single target survives, to accomplish which he should distribute the defenders as evenly as possible over the attackers. Let r be the remainder when m is divided by n :

$$(4-1) \quad m = kn + r, \quad \text{where} \quad 0 \leq r < n.$$

When the defenders are distributed as evenly as possible, $(n - r)$ attackers are assigned k interceptors, r are assigned $k + 1$, and the probability that the target survives is

$$(4-2) \quad Q(m, n) \equiv [1 - p(1 - \rho)^k]^{n-r} [1 - p(1 - \rho)^{k+1}]^r.$$

For example, suppose $p = .8$, $\rho = .5$, $m = 7$, and $n = 3$, so that $k = 2$ and $r = 1$. Each of the 2 attackers that are assigned 2 interceptors will kill the target with probability $p(1 - \rho)^2 = .2$, and the target will therefore survive both attackers with probability $.8^2 =$

.64, which is the first [] factor in (4-2). The second is $(1 - .8(5)^3)^1 = .9$, so $Q(7, 3) = .576$.

The target survival probability $Q(m, n)$ can be approximated by permitting non-integer allocations of interceptors, m/n to each attacker:

$$(4-3) \quad Q(m, n) = (1 - p(1 - \rho)^{m/n})^n$$

Equation (4-3) approximates $Q(7, 3)$ in the previous example by .595; (4-3) will in all cases be at least as large as (4-2).

The conclusion that interceptors should be evenly distributed also holds if the defender's goal is to destroy as many attackers as possible, on the average. This might be a reasonable goal if p were unknown, or if the attackers were not all aimed at the same target. Let $A(m, n)$ be the average number of attackers out of n that survive the m interceptors. Exercise 11 is to find an expression for $A(m, n)$ and record it in the space provided below:

$$(4-4) \quad A(m, n) =$$

Suppose now that the attackers arrive one at a time, and let m_i be the number of interceptors allocated to the i^{th} attacker. We have just concluded that the m_i should be as equal as possible (a "flat" defense), but the reader may have intuitive feelings that a "tapered" defense would be more desirable; i.e., that m_1 should be larger than m_2 , etc. There are a variety of reasons why a tapered defense might actually be a good idea, the most important of which is the possibility that the total number of attackers might be unknown. Sections 4.2, 4.3, and 4.4 deal with three distinct versions of the problem where n is unknown. Even in the case where the objective is to shoot down as many attackers as possible, a tapered defense might still be advisable if the defensive system were part of the target; i.e., if no further interceptors could be launched after a target kill. If the objective is to maximize the target survival probability, however, the best defense

is flat as long as n is known. This is true even if the attack is sequential and if the defensive system is part of the target.

We return now to the case where all attackers appear simultaneously, and where the object of the defense is to maximize the target's survival probability, but we suppose that the defender has the time and information required to implement a J stage shoot-look-shoot policy; i.e., the defender can shoot at the attackers, then shoot at the survivors, etc., until either no attackers remain or only one stage remains, in which case all remaining interceptors should be distributed evenly over whatever attackers are still alive. In each stage, the defender can use as many interceptors as he likes. If J is very large, the defender can safely adopt the strategy of firing one interceptor at each surviving attacker at each stage until either no attackers or no interceptors remain, but this strategy is not optimal if J is small (the analysis culminating in (4-2) corresponds to the special case $J=1$, in which the defender may very well fire more than one interceptor per attacker).

The problem of computing the optimal firing policy at each stage is non-trivial, even if one recognizes at the outset that interceptors should still be distributed evenly over attackers at each stage. The difficulty is due to the fact that it is not obvious how many interceptors should be used per stage, except in the last stage, and that in any case the best number to use probably depends on how many attackers survive, which is random. In other words, the form of the optimal policy is not "use 5 in the first stage, 3 in the second, ...", but rather, "use 5 in the first stage, 4 in the second if 4 attackers survive the first, six in the second if 3 attackers survive the first, ..."; i.e., the actual optimal policy must involve a great many statements that are conditional on the results of earlier stages. There are a great many policies of the latter form — so many that one would not even consider solving a non-trivial problem by examining all of them, even on a computer. Nonetheless, the optimal policy is not difficult to determine. The technique required is

Dynamic Programming (DP), a recursive method that involves the idea of “state”, an idea that is fully as important in Operations Research as in Physics.

Dynamic Programming problems usually involve an evolving process of some sort, with the definition of “state” being whatever information about the past is sufficient for purposes of taking action in the future. Correct identification of the state is a crucial part of Dynamic Programming problem formulation. It may help to imagine a “change of command” in the middle of the process, with the state of the process being whatever information the old commander should transfer to the new one. In the problem under consideration, the state is (j, m, n) , where the three variables are the number of stages left, the number of interceptors left, and the number of attackers still surviving, respectively. The past may have much more detail than that, but all such detail is irrelevant for purposes of future action — three numbers suffice.

Given the state, there are two more steps to be taken in the successful formulation of a DP. The first step is simple: write down what the objective function of the new commander should be, both as a mathematical function of the state and in English. The English part is essential. For our firing problem,

$$(4-5) \quad F(j, m, n) = \text{“the maximum possible probability of surviving } n \text{ attackers if } m \text{ interceptors and } j \text{ stages remain”}.$$

Let M and N be the total number of interceptors and attackers, respectively. These were previously called m and n , but m and n are now being used as dummy variables. The number that we seek is $F(J,M,N)$, together with the associated firing policy. The function $F(1,\cdot,\cdot)$ is already known, being given by (4-2). The idea is to use $F(1,\cdot,\cdot)$ to compute $F(2,\cdot,\cdot)$, then $F(2,\cdot,\cdot)$ to compute $F(3,\cdot,\cdot)$, etc., until finally $F(J,\cdot,\cdot)$ is obtained, after which $F(J,M,N)$ is a special case. The second step in a DP formulation is the construction of a recursive formula that accomplishes this. For our firing problem,

$$(4-6) \quad F(j+1, m, n) = \max_{0 \leq u \leq m} \{ E(F(j, m-u, X)) \},$$

where X is the random number of attackers that survive the current stage. The crucial thing about (4-6) is that $F(j, \cdot, \cdot)$ appears on the right hand side and $F(j+1, \cdot, \cdot)$ on the left. (4-6) must be evaluated for $1 \leq j < J$, $0 \leq m \leq M$, and $0 \leq n \leq N$, a total of $(J-1)(M+1)(N+1)$ times. If u is the number of interceptors utilized in state $(j+1, m, n)$, then the state will be $(j, m-u, X)$ when the next decision is made; the probability distribution of X depends on u , and the fact that X is random requires the expectation operation. The amount of computation required is considerable, but quite feasible on a computer, and nowhere near as much as would be required to examine all possible firing strategies. The optimal number of interceptors to fire is simply the maximizing value of u obtained in the process of computing (4-6); call it $u^*(j+1, m, n)$. By recording the function $u^*(\cdot, \cdot, \cdot)$, the defense is prepared for all possible eventualities.

There are three applications of DP. In these notes, the most difficult of which is probably the one just discussed. The reader who is unfamiliar with DP may prefer to begin with the more elementary applications in Secs. 4.2 and 4.3 (especially 4.3). This application is continued in exercise 12.

4.2 Bayesian defense

In this section a stockpile of M interceptors, each of which has kill probability ρ , is available for the defense of a single target against a sequence of attackers, each of which has kill probability p if not intercepted. The total number of attackers (A) is unknown, but is assumed to be no larger than some number N . For example, N attackers might be committed to the attack, with an unknown number of them either malfunctioning or being destroyed by other defensive systems. It is assumed that enough is known about the process to construct the probability law for the random variable A , so the quantities $P(A > i)$ are assumed known for $i = 0, \dots, N$, with $P(A > 0) = 1$. Let m_i be the number of interceptors allocated to the i^{th} attacker. The object is to determine the firing schedule

m_1, \dots, m_N that maximizes the probability of surviving all A attackers subject to the constraint that $m_1 + \dots + m_N = M$. Section 4.1 would apply if A were known, but A is random.

We will solve the problem using Dynamic Programming. The state of the process is (m, i) , where m is the number of interceptors remaining and i is the number of attackers that have already arrived. The decision required is to determine the number of interceptors u to use against the next attacker. The state must include i because if $i = N - 1$, for example, then it is clear that m is the best choice for u , whereas it may be wise to make $u < m$ if $i = 1$. The objective function is

(4-7) $F(m, i)$ = “the maximum probability of surviving all future attackers if i have already arrived (without killing the target), and if m interceptors remain”

To develop the recursive formula for $F(m, i)$, we must first recognize that there may not be any future attackers at all, in which case survival is certain. The probability that there will be at least one more attacker, given that i attackers have already arrived, is $Q_i \equiv P(A > i + 1 | A > i) = P(A > i + 1)/P(A > i)$. If there is at least one more attacker, and if the next attacker does not destroy the target, then the next state will be $(m - u, i + 1)$. The desired recursion is therefore

$$(4-8) \quad F(m, i) = 1 - Q_i + Q_i \max_{0 \leq u \leq m} \left\{ [1 - p(1 - \rho)^u] F(m - u, i + 1) \right\}.$$

It is clear that $F(\cdot, N) = 1$, since survival is certain if all attackers have already arrived. (4-8) can therefore be used to compute $F(\cdot, N - 1)$, then $F(\cdot, N - 2)$, etc., until finally $F(\cdot, 0)$ is obtained. In the process of doing the computations, the optimal allocation of interceptors can be recorded as $u^*(m, i)$, and this determines the optimal firing schedule. The number of interceptors to be allocated to the first attacker is $m_1 = u^*(M, 0)$, and then $m_2 = u^*(M - m_1, 1)$, etc. See exercise 13.

4.3 The maximum cost defense

We assume here, as in Section 4.2, that the number of attackers is unknown, and that a firing schedule for the defensive interceptors must nonetheless be set up for use against a sequence of attackers. However, no probability distribution is given for the total number of attackers. Instead, the defense takes the point of view that any target defended by a finite stockpile of interceptors can be killed if sufficiently many attackers are committed, and that the proper goal is therefore to maximize the cost (measured in attackers) of killing the target. If this number turns out to be so large that the attack does not take place, then so much the better, but in any case the defensive goal is to make the target as hard to kill as possible. The attacker is assumed to have a shoot-look-shoot capability.

The objective of maximizing the average number of attackers required to kill the target can be accomplished using Dynamic Programming. The state of the process is simply the number of interceptors m remaining, and the objective function is

$$(4-9) \quad c(m) = \text{"the average number of additional attackers required to kill the target if } m \text{ interceptors remain".}$$

Suppose u interceptors are allocated to the next attacker. The probability that the next attacker kills the target is then $p(1 - \rho)^u$, where ρ and p are the kill probabilities of attackers and interceptors, respectively. If the next attacker fails to kill the target, then the next state will be $m - u$. Therefore,

$$(4-10) \quad c(m) = 1 + \max_{0 \leq u \leq m} \left\{ (1 - p(1 - \rho)^u) c(m - u) \right\}.$$

If $m = 0$, (4-10) is the equation $c(0) = 1 + (1 - p)c(0)$, which has the solution $c(0) = 1/p$. This is the average number of attackers required to kill an undefended target. For $m > 0$, the option $u = 0$ can safely be ignored, since at least one interceptor should be used in any case. (4-10) can now be used to determine $c(1)$, then $c(2)$, etc., recording the maximizing

value of u at each stage (call it $u^*(m)$). For example, suppose $p = .8$ and $\rho = .5$. Then $c(0) = 1.25$, and (the maximizing element is underlined)

$$c(1) = 1 + .6c(0) = 1.75, \quad \text{and } u^*(1) = 1$$

$$c(2) = 1 + \max\{\underline{.6c(1)}, .8c(0)\} = 2.05, \quad \text{and } u^*(2) = 1$$

$$c(3) = 1 + \max\{.6c(2), \underline{.8c(1)}, .9c(0)\} = 2.40, \quad \text{and } u^*(3) = 2$$

Continuing in this manner, we find that $c(m) = 1.25, 1.75, 2.05, 2.40, 2.64, 2.92, 3.16, 3.38, 3.63, 3.84, 4.04, 4.27$ for $m = 0, 1, \dots, 11$, and also $u^*(m) = 0, 1, 1, 2, 2, 2, 3, 3, 3, 3, 3, 3$. If 11 interceptors remain, 3 should be used against the first attacker, then $u^*(8) = 3$ should be used against the second, $u^*(5) = 2$ against the third, $u^*(3) = 2$ against the fourth, and $u^*(1) = 1$ against the fifth. The sixth and subsequent attackers would not be opposed, assuming that six or more were actually required to kill the target.

The function $c(\bullet)$ is not tactically necessary, since the firing schedule is implicit in the function $u^*(\bullet)$. One might, however, use $c(m)$ as a measure of effectiveness for making a quantity vs. quality decision (see exercise 15).

4.4 Prim-Read defense

The assumptions in this section are the same as in Section 4.3, except that the attacker is no longer assumed to have a shoot-look-shoot capability. The attackers still arrive sequentially, but a certain number (say n) out of a large stockpile must be irrevocably committed to the target. Let $p(n)$ be the probability that the target is killed by one of n attackers, and let $\lambda = \max_{n>1} p(n)/n$. λ is the largest possible kill probability per attacker.

The objective of a Prim-Read defense is to make λ as small as possible, the idea being to prevent “cheap kills”. The idea was first proposed as a method for defending targets with ABM’s against ICBM attack.

The problem of minimizing the defensive stockpile required to achieve a given λ turns out to be much easier than the problem of minimizing λ for a given stockpile; so much so that a problem of the latter type is most easily solved by guessing values for λ .

until the calculated stockpile is whatever happens to be available. This technique is illustrated below. Let m_i be the number of interceptors allocated to the i^{th} attacker. Making the usual independence assumptions, and letting p and ρ be the kill probabilities of attackers and interceptors,

$$(4-11) \quad p(n) = 1 - \prod_{i=1}^n (1 - p(1 - \rho)^{m_i}); \quad n \geq 1,$$

and the central problem is to minimize $\sum_{i=1}^{\infty} m_i$ subject to the constraints that $p(n) \leq \lambda n$

for all $n \geq 1$. Suppose, for example, that $p = .8$, $\rho = .5$, and that there are $m = 11$ interceptors available. Our initial guess is that 11 interceptors should be sufficient to guarantee that the kill probability per attacker need not exceed (say) $\lambda = .3$. We now consider the problem of minimizing the number of interceptors required to guarantee that the maximum kill probability per attacker does not exceed .3, hoping that the answer is 11. From (4-11), $p(1) = p(1 - \rho)^{m_1}$. Since $p(1)$ must not exceed .3, the smallest possible value for m_1 is 2, so we take $m_1 = 2$. From (4-11), we therefore have $p(2) = 1 - .8(1 - p(1 - \rho)^{m_2})$. The smallest value of m_2 for which $p(2) < .6$ is 1, so we take $m_2 = 1$. From (4-11), $p(3) = 1 - (.8)(.6)(1 - p(1 - \rho)^{m_3})$, and the smallest value of m_3 for which $p(3) < .9$ is 1 ($p(3) = .904$ when $m_3 = 0$, which is just barely too large), so we take $m_3 = 1$. Since $.3n > 1$ for $n \geq 4$, $m_i = 0$ for $n \geq 4$. The total number of interceptors required to guarantee that the kill probability per attacker does not exceed .3 is therefore $2 + 1 + 1 = 4$. Eleven interceptors are evidently sufficient for a smaller value of λ . The next step is to guess a smaller value (see exercise 17) and repeat the above calculations. The calculations are easy because the product in (4-11) can be formed sequentially, with the first $(n - 1)$ factors being known when m_n is being determined. The process could be easily automated in a spreadsheet. The easiness of the calculations makes up for the fact that they must typically be repeated several times.

Although a Prim-Read defense can certainly be constructed for a single target, the technique is more naturally applied to a group of several targets, using the same value of λ for every target in the group. If the targets differ from each other, one simply replaces $p(n)$ with the function $v(n) \equiv$ “avg value killed by n attackers”. An implicit assumption in setting up such a defense is that the attacker can determine the defensive firing schedule before making his own allocations. There may be good physical reasons for assuming this, but it may also be true that the attacker has just as much trouble ascertaining defensive allocations as vice versa. In the latter case, a Prim-Read defense is probably a mistake. The Prim-Read defense of several identical targets would treat all targets equally, for example, whereas the best defense may be to abandon half of the targets in order to construct a strong defense of the remainder. The natural way to formulate such a problem would be as a two person zero sum game, but doing so is beyond the scope of these notes.


Exercises

- 1) Suppose $D(r) = 1 - r$ if $r < 1$; 0 if $r > 1$. What is the lethal area?

Ans. $a = \pi/3$

- 2) Plot $D(r)$ for the target illustrated below, assuming that the weapon must hit the shaded area and that the impact point is (r, θ) with θ uniformly random in $[0, 2\pi]$. Show that the lethal area is equal to the area of the target.

Ans. $D(r)$ is a step function, $a = 2.5\pi$



- 3) Show that (2-4) produces $\pi E(R^2)$ for lethal area, where $E(R^2)$ is computed using (2-6). Hint: use integration by parts.
- 4) Derive (2-7).
- 5) When aiming errors are basically angular, the miss distances should increase with range. Suppose several independent shots are taken at a target, with $\sigma_i = .1r_i$, where r_i is the i^{th} range, and that the cookie cutter lethal radius is 1. If the successive ranges are 10, 11, 12, etc., compute P_K for the first shot, the first five shots (as a group), and the first 10 shots.

Ans. ($P_K(1) = .39$, $P_K(5) = .84$, $P_K(10) = .93$).

- 6) Justify the $\exp(-d)$ formula that was used in deriving (3-7). Hint: Argue that the number of times any given point is covered is a Poisson random variable with mean d .
- 7) An aircraft attempts to kill a tank as follows: It first drops a canister of “stickers” in the hope that one will hit the tank and activate. If a sticker activates, it can guide a projectile to the tank. The canister opens and scatters 1000 stickers, with the amount of scatter being under the control of the designer. The exposed area of the tank is 100 yd^2 . The aircraft makes a 2-dimensional error with standard deviations (100 yards, 300 yards) in dropping the canister. What is the probability that a sticker hits the tank, assuming a well-designed canister? If the tank is longer than it is wide, does the direction of the aircraft’s approach matter?

Ans. (.275, no)

- 8) Suppose you are given 16 detection devices, each of which is guaranteed to detect a target if and only if the relative distance is either less than 4 miles or between 30 and 33 miles (the “convergence zone” phenomenon in the ocean might be one explanation for such an assumption). The devices can be placed in any pattern whatever, and the object is to detect a target whose location relative to some known

point is circular normal with standard deviation 30 miles in each direction. There are no dispersion errors.

- a) Estimate p_K .
- b) Make up a pattern and test it by writing a 5000-replication computer simulation.

Ans. The lethal area is $\pi(4^2 + 33^2 - 30^2) = 205\pi$, so $z = (16)(205)/1800 = 1.82$.

Given that the shape of the lethal area makes considerable overlap inevitable even in the absence of dispersion, a confetti approximation is natural. The $\sqrt{-}$ formula produces $P_K \approx .57$. This example has been the result of considerable experimentation, with the best pattern as of this writing having a detection probability of .64.

- 9) Suppose 10 cookie cutter shots are available, with the lethal radius being 30 ft. for each. Estimate P_K for the area target and errors considered in Sec. 3.3, assuming that
 - a) the wind error is dispersion
 - b) the wind error is bias

Ans. Using the $\sqrt{-}$ formula in both cases, the expected fraction of the target killed with an optimal pattern would be approximately .316 in case a), or .275 in case b).

- 10) If the SAS procedure (3-12) is used for aim adjustment, then the miss probability with the i^{th} shot is $q_i = \exp[-(R^2/2\sigma^2)(i-1)/i]$; $i > 1$. Use this fact along with (3-16) to compute $E(N)$ when $(R^2/2\sigma^2) = \ln(2)$. It will be necessary to write a computer program. Note that only 2 shots would be required, on the average, if there were no bias error, since each shot would have a kill probability of .5.

Ans. $E(N) = 3.76$.

- 11) Do the exercise described in Sec. 4.1. Hint: $A(7,3) = 5/8$ when $p = .5$.

12) The Dynamic Programming sheet of Firing.xls implements (4-6) for up to 7 attackers, 12 defenders, and 3 stages, with the p and ρ parameters of Sec. 4.1 being inputs (there is a separate p for each stage). The optimized function $F(j, m, n)$ is shown in yellow cells, and the optimizing allocation of defenders to targets $u^*(j, m, n)$ is shown in green cells. Solve several problems and state whether you think each of the following propositions is true or false when all stages have the same p :

- a) For fixed (j, n) , $u^*(j, m, n)$ is a nondecreasing function of m .
- b) All attackers are treated equally; that is $u^*(j, m, n)$ is always an integer multiple of n .
- c) $u^*(j, m, n) \geq \min(m, n)$.
- d) For fixed (j, m) , $u^*(j, m, n)$ is a nondecreasing function of n .

Also, give a qualitative description of what happens if the last stage has a higher p than the other 2, as might happen if defenders were more effective at close range.

13) Suppose that the number of attackers A is random, with $P(A = i) = .1, .3, .4, .2$ for $i = 0, 1, 2, 3$. The attackers, whatever their number, arrive sequentially, so that interceptors must be allocated to each attacker as it appears. Each interceptor kills its target with probability .5. Using (4-8), determine the firing schedule that maximizes the probability of destroying all attackers if there are 4 interceptors in total.

Ans. $m_1 = 2, m_2 = 2, m_3 = 0$. The probability of destroying all attackers is $F(4,0) = .55$.

14) Modify (4-7) and (4-8) to reflect the goal of shooting down the largest number of attackers on the average, and determine the optimal firing schedule using the parameters given in problem 13

- a) assuming that the defense is invulnerable;
- b) assuming that no further attackers can be shot down if an attacker succeeds in killing the target.

15) Sec. 4.3 includes an example where $\rho = .5$, $p = .8$, and $m = 11$. Would the defender prefer 4 perfect ($\rho = 1$) interceptors to 11 imperfect ($\rho = .5$) ones

- a) using the measure of effectiveness $c(m)$?
- b) using the measure of effectiveness ρm ?

Ans. Perfect in a), imperfect in b). Note the high emphasis that $c(m)$ places on quality.

16) Using a modified version of (4-10), compute $c(8)$ under the assumptions that $p = .8$, $\rho = .5$, and that at most two interceptors can be allocated to any attacker.

Ans. $c(8) = 3.49$.

17) Continue the sample analysis begun in Sec. 4.4 by next guessing $\lambda = .15$. You should find that the required firing schedule utilizes 11 interceptors. Compare your calculations with the firing schedule obtained in Sec. 4.3 using the same parameters.

REFERENCES

- [1] Eckler, A.R. and Burr, S.A., 1972. "Mathematical Models of Target Coverage and Missile Allocation", Military Operations Research Society.
- [2] "Mathematical Background and Programming Aids for the Physical Vulnerability System", DI-550-27-74, Defense Intelligence Agency (1974).
- [3] Tsipis, K., 1974. "The Calculus of Nuclear Counterforce", Technology Review, pp. 34-47.
- [4] Morse, P.M. and Kimball, G.F., 1950. "Methods of Operations Research".

ADDITIONAL READING

- [1] Svesnikov, A.A., 1948. "Collection of Articles on the Theory of Firing", (A.N. Kolmogorov, ed.). Translation T-14, The RAND Corporation.
- [2] Weiss, H.K., 1953. "Methods for Computing the Effectiveness of Area Weapons", Ballistic Research Laboratory Report 879.
- [3] "Probability-of-Damage Problems of Frequent Occurrence", Operations Evaluation Group Study No. 626 (1959).
- [4] Swinson, G.E., Randolph, P.H., Dunn, B.J. and Walker, M.E., 1971. "A Model for Allocating Interceptors From Overlapping Batteries: A Method of Dynamic Programming", Operations Research, vol. 19, pp. 182-193.
- [5] Soland, R.M., 1973. "Optimal Defensive Missile Allocation: A Discrete Min-Max Problem", Operations Research, vol. 21, pp. 590-596.
- [6] Shumate, K.C. and Howard, G.T., 1974. "A Proportional Defense Model", Naval Research Logistics Quarterly, vol. 21, pp. 69-78.
- [7] Sverdlov, E.E., 1981. "Optimal Allocation of Tactical Missiles Between Valued Targets and Defense Targets", Ph.D. Thesis, US Naval Postgraduate School, Monterey.
- [8] Grotte, J.H., 1982. "An Optimizing Nuclear Exchange Model for the Analysis of Nuclear War and Deterrence", Operations Research, vol. 30, pp. 428-445.